arXiv:1612.01234v1 [cs.CV] 5 Dec 2016

Multi-way Particle Swarm Fusion

Chen Liu*[1], Hang Yan*[1], Pushmeet Kohli[2], and Yasutaka Furukawa[1]

[1] {chenliu,yanhang,furukawa}@wustl.edu
[2] pkohli@microsoft.com

Abstract. This paper proposes a novel MAP inference framework for Markov Random Field (MRF) in parallel computing environments. The inference framework, dubbed Swarm Fusion, is a natural generalization of the Fusion Move method. Every thread (in a case of multi-threading environments) maintains and updates a solution. At each iteration, a thread can generate arbitrary number of solution proposals and take arbitrary number of concurrent solutions from the other threads to perform multi-way fusion in updating its solution. The framework is general, making popular existing inference techniques such as alpha-expansion, fusion move, parallel alpha-expansion, and hierarchical fusion, its special cases. We have evaluated the effectiveness of our approach against competing methods on three problems of varying difficulties, in particular, the stereo, the optical flow, and the layered depthmap estimation problems. [1]

Keywords: MRF; Fusion Move; Particle Swarm Optimization

1 Introduction

Parallel computation has changed the field of computing. In the 90s, most processors had single cores. In 2016, processors have often 4 cores, or even 8. Cluster computing further expands the potential of parallel computation, where one can easily launch a processing job using hundreds or even thousands of computational nodes in a cloud. In the recent work on the AI program playing the ancient Chinese board game of *Go*, parallelization plays a key role in the Monte-Carlo tree search [1].

Parallel computation offers tremendous potential for Computer Vision. As image sensing technologies have gone through revolutions, we are in ever growing demands in solving very large problems. One may need to apply image denoising to 50 Megapixel images from latest digital SLRs (e.g., Canon EOS 5DS), stitch thousands of images to generate gigapixel panoramas [2], or solve volumetric reconstruction and segmentation problems over a billion ($= 1024^3$) voxels [3]. Markov Random Field (MRF) has been a very successful framework to solve these problems in Computer Vision. However, state-of-the-art algorithms for MRF inference are still inherently sequential. Take a Fusion Move method

*indicates equal contribution.

[1]Project page: http://www.cse.wustl.edu/∼chenliu/swarm-fusion.html

(FM) [4–6] for example, which has been one of the most effective techniques for MRF inference. It sequentially improves solution by fusing the current solution with a solution proposal. It has been successfully applied to many problems such as optical flow, stereo, image inpainting, or image segmentation [5].

Unleashing the power of parallel computation for effective MRF inference would then bring fundamental contributions to Computer Vision. Currently, FM suffers from a few vital limitations due to its sequential nature. First, standard FM allows only two options per variable in each fusion, either the current solution or a proposal [5]. Second, only a single proposal generation scheme is used in each fusion step. [1] Our approach, dubbed *Swarm Fusion* method (SF), makes a few key distinctions from existing approaches: 1) Multiple threads (or computing nodes) simultaneously keep and improve solutions; and 2) Each fusion in each thread can generate arbitrary number of solution proposals and use arbitrary number of concurrent solutions in the other threads, to be fused with the current solution.

We have evaluated the effectiveness of our approach over three problems in Computer Vision, specifically, stereo, optical flow, and layered depthmap estimation. Our idea is extremely simple and the new inference framework can be integrated into existing system with minimal coding. We believe that this paper would have immediate impact on numerous Computer Vision researchers or engineers, currently solving MRF problems with conventional methods.

2 Related work

MRF inference has been a very active field in Computer Vision with extensive literature. We refer the readers to survey articles for comprehensive reviews [8,9], and here focus our description on closely related topics.

Parallel Alpha-Expansion

Lempitsky et al. [5] introduces parallel computation to the alpha-expansion technique, where multiple threads simultaneously fuse mutually exclusive sets of labels. Kumar et al. [10], Delong et al. [11], and Veksler et al. [12] investigated hierarchical approachs, where labels can be simultaneously fused from the bottom to the top in a tree of labels. Instead of taking a hierarchical approach, Batra et al. [13] adaptively computed an effective sequence of labels to explore. This technique can be combined with parallel alpha-expansion techniques to obtain further speed-up. Strictly speaking, these approaches are not in the family of Fusion Move methods (FM), because they only consider constant label proposals. Our approach is a generalization of FM.

Parallel MAP inference

[1] Recently, an extension of FM was introduced for layered depthmap estimation [7], where a solution subspace, instead of a single solution, is proposed and fused with the current solution. However, this approach is also limited to the use of one proposal generation scheme in each fusion.

The core MAP inference itself can be parallelized. Strandmark et al. [14] parallelized graph-cuts. Message passing algorithms are friendly to GPU implementation and can exploit the power of parallel computation. While state-of-the-art optimization libraries are often freely available for non-commercial purposes, most companies have to develop and maintain in-house implementation of these algorithms. The core optimization libraries are very complex and their modifications require significant engineering investments. In contrast, our idea is extremely simple and easily reproducible by standard engineers.

Fusion Move methods

FM was first introduced by Lempitsky et al. [4] in solving the optical flow problem. FM has been effectively used to solve other challenging problems in Computer Vision such as stereo with second order smoothness priors [6], stereo with parameteric surface fitting and segmentation (i.e. Surface Stereo) [15], and multicut partitioning [16]. FM has two main advantages over other general inference techniques [17, 18]. First, FM allows us to exploit domain-specific knowledge by customizing proposal generation schemes. Second, FM can handle problems with very large label spaces (and even real-values variables), because the core optimization solves a sequence of binary decision problems. In contrast, methods like message passing algorithms need to maintain messages and beliefs for the entire label space all the time. Although conceptually straightforward, we are not aware of *Parallel Fusion Move (PFM)* algorithms that fuse solution proposals, as opposed to labels, in parallel. This paper seeks to fully unleash the power of parallel computation based on FM in the most general setting.

Evolutionary algorithms and Particle Swarm Optimization

Genetic algorithms (GA) [19] and Particle Swarm Optimization (PSO) [20] maintain multiple solutions and improve them over time. GA or PSO has been used to produce great empirical results, e.g. in hand tracking [21]. At high level, our strategy is similar in spirit. However, GA or PSO rather arbitrarily copies parts of the solutions or makes random movements in each step (i.e., limited theoretical justification). Our approach directly optimizes the objective function to improve solutions.

3 Multi-way particle swarm fusion

Multi-way Particle Swarm Fusion is a natural extension of the Fusion Move method (FM). We call our method Swarm Fusion (SF) in short. Let us take multi-threading environment to explain our idea, while the technique is also applicable to other parallel programming model such as *MapReduce* in cloud computing.

Assuming we have N threads $\{T_i | i = 1, 2, \cdots, N\}$, each thread T_i maintains and updates a solution S_i in parallel. SF has 1) a *proposal generator* for each thread which picks arbitrary number of proposal generation schemes and generates proposals, and 2) a *solution pool*, from which a thread picks arbitrary

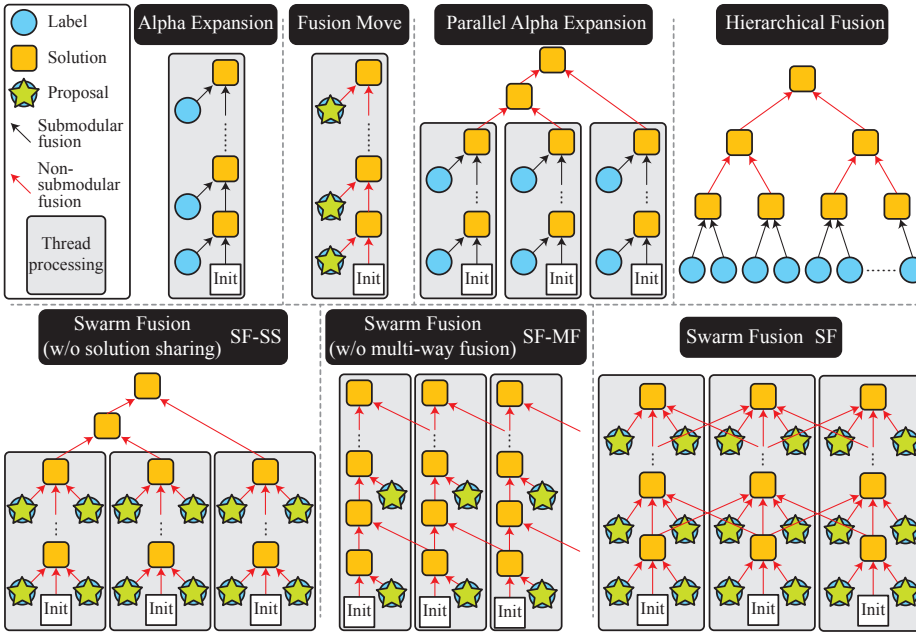

Fig. 1: Swarm Fusion (SF) architecture and its relationships to existing methods. The bottom right example shows the general SF architecture, where each thread takes arbitrary number of solution proposals and concurrent solutions for fusion. The framework is flexible and can realize other data processing architectures depending on the parameters (e.g., the left two examples in the bottom row). It is easy to verify that existing popular MRF inference methods such as Alpha Expansion [22], Fusion Move [5], Parallel Alpha Expansion [12], and Hierarchical Fusion [11], are all special cases of SF.

number of intermediate solutions generated by the others. In our base configuration, the solution pool remembers N best solutions, one from each thread.

SF has two main parameters α_i, β_i (for each thread T_i), determining its behaviors: In each fusion step, a thread generates α_i solution proposals using its proposal generator, and collects β_i solutions from the solution pool, based on a user-defined strategy or at random to be simple. The values of α_i and β_i can vary per iteration for flexibility. The thread then fuse all these proposals and/or solutions to find a solution with lower energy state and update the solution pool accordingly.

Swarm Fusion framework is very flexible and yields various data processing architectures as shown in Fig. 1. The bottom right architecture is the most general one, in which threads conduct multi-way fusion of their current solution, proposals from their own proposal generators and/or concurrent solutions. For general non-submodular energy, we use TRW-S [17] for inference. However, if one knows that a certain fusion step is a binary fusion with submodular energy,

Algorithm 1 Swarm Fusion method

procedure (α, β)
  $\mathcal{S}_{pool} \leftarrow \emptyset$ //Solution pool
  for each thread T_i **do**
    Initialize its solution S_i
  end for

  for each thread T_i in parallel till convergence **do**
    Generate α_i solution proposals $\mathcal{P}$
    Pick β_i solutions $\mathcal{S} \subset \mathcal{S}_{\text{pool}}$
    $S_i \leftarrow \text{Fuse}(S_i, \mathcal{P}, \mathcal{S})$
    Replaces the solution in $\mathcal{S}_{pool}$ with S_i
  end for
end procedure

one can use alpha-expansion [22]. QPBO [23] can be used to perform binary fusion with non-submodular energy. Note that the threads appear synchronized in the figure only for illustration purpose. In practice, all the threads run asynchronously with a (read-write) lock on the data in the solution pool (See Algorithm 1).

Relationships to existing methods

It is easy to verify that Alpha-Expansion (AE) [22], Fusion Move (FM) [5], Parallel Alpha Expansion (PAE) [5], and Hierarchical Fusion (HF) [11, 12] are all special cases of the Swarm Fusion method (SF). AE can be realized by setting $(\alpha = 1, \beta = 0)$ and restricting the proposals to be constant labels with a single thread. The same goes for FM, this time, without the restriction on the proposal generation scheme. PAE is realized by setting $(\alpha = 1, \beta = 0)$ with multiple threads, again with a restriction on the proposal generation scheme (the last sequential fusion in PAE is realized by $(\alpha = 0, \beta = 1)$ with a single thread). HF has a slightly different data processing model, without strong ties between threads and data, but can be realized by setting $(\alpha = 2, \beta = 0)$ at the bottom level and $(\alpha = 0, \beta = 2)$ at the remaining levels, while allowing S_i not to be used in the fusion steps of T_i.

4 Swarm Fusion instantiation

We compare SF against competing approaches over three problems in Computer Vision, specifically, stereo, optical flow, and layered depthmap estimation (see Fig. 2).

4.1 Swarm Fusion stereo

We start with a simple depthmap stereo problem with standard unary and pairwise terms. We employ submodular pairwise terms to make this stereo represent

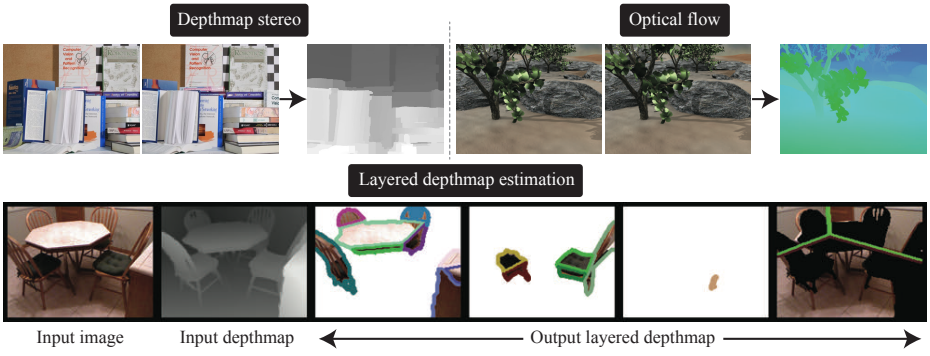

Fig. 2: We compare our Swarm Fusion method against competing approaches on the depthmap stereo [24], the optical flow [25] and the layered depthmap estimation [7] problem. In the layered depthmap problem, the input is a RGBD image, and the output is multiple layers of depthmaps. Each layer is a piecewise smooth parametric surface model.

relatively "easy" MRF inference problem. The unary terms are computed as the average robust photoconsistancy score [6] between the reference image and the others inside a 7×7 pixels window. The pairwise terms are simple truncated absolute label difference with maximum label difference $\sigma_s = 4$. The total energy is defined by the sum of the two, while scaling the pairwise terms by a factor of 0.005. For simplicity we do not enforce the visibility constraint.

Competing methods

For simple stereo problems with submodular energy as ours, the sophistication of photometric consistency function [26] makes unary terms highly informative, where efficient inference algorithms such as graph-cuts exist. Therefore, we have chosen algorithms based on Alpha-Expansion, namely single thread Alpha Expansion(AE), Parallel Alpha Expansion(PAE) [5] and Hierarchical Fusion(HF) [12] to be competing methods. For HF, we use Alpha-Expansion at the leaf node of the label tree and QPBO in the other cases.

Swarm Fusion architectures

The three swarm architectures in Fig. 1 have been evaluated: SF-MF (SF without multi-way fusion), SF-SS (SF without solution sharing), and the standard SF. SF-MF implies $\alpha + \beta = 1$, where each thread repeats fusing a solution proposal $(\alpha = 1, \beta = 0)$ for four iterations by Graph-cuts and fusing a concurrent solution $(\alpha = 0, \beta = 1)$ for one iteration by QPBO. In the later case, a thread randomly chooses one solution from the solution pool for fusion. SF-SS implies $\beta = 0$, where α is the free parameter and set to 4. In this case one thread will fuse 4 labels, together with current solution in that thread by TRW-S in each iteration and never exchanges solutions with other threads. We perform a multi-way fusion of solutions from all the threads at the end to obtain a final solution (similar

to PAE). For standard SF architecture, we have used ($\alpha = 4, \beta = 1$). To make the comparison simple, we restrict our solution proposals to be constant-label proposals.

4.2 Swarm Fusion optical flow

Fusion Move was first introduced by Lempitsky et al. [4] to solve the optical flow problem. We copy their problem setting and use images from the Middlebury optical flow benchmark [25]. We share similar proposal generation schemes with Lempitsky et al [4] with some modifications. [1]

Competing methods

Fusion Move method in Lempitsky's paper is the first natural contender. While they did not consider parallel implementation, it is straightforward to combine the idea of Parallel Alpha Expansion and Fusion Move. Therefore, the second competing method is "Parallel Fusion Move"(PFM), which is equivalent to Parallel Alpha Expansion with constant label solutions replaced by solution proposals. One problem of PFM is that infinite number of solution proposals can be generated in their algorithm, and we do not know when to stop and perform the final sequential fusion (See Parallel Alpha Expansion architecture in Fig. 1). In our experiments, we manually picked time limits to initiate the final fusion to make the comparisons fair. The last contender is the mix of the Hierarchical Fusion and the Fusion Move methods, dubbed "Hierarchical Fusion Move"(HFM), where they start from solution proposals as opposed to constant labels. One problem of HFM is that we need to generate all the proposals first to build the fusion tree. This undermines the power of fusion move that can adaptively generate proposals based on the current solution. In our experiments, we have manually generate 250 proposals at the beginning. The fusions are binary in these methods and we have used QPBO.

Swarm Fusion architectures

The three swarm architectures in Fig. 1 (SF-MF, SF-SS, SF) have been evaluated against the competing methods. For SF-MF, each thread repeats generating solution proposals ($\alpha = 1, \beta = 0$) for four iterations and fuses with one solution from others ($\alpha = 0, \beta = 1$) for one iteration. This pattern is repeated. For SF-SS, each thread generates three solution proposals for fusion in each iteration ($\alpha = 3$, $\beta = 0$). For SF, we repeat four iterations of ($\alpha = 3, \beta = 0$) and one iteration of ($\alpha = 0, \beta = 3$). We have used TRW-S for multi-way fusion and QPBO for binary fusion.

[1] First, we use more recent Farneback algorithm and change the level of pyramids from 1 to 5, then use either 3, 5 or 7 for parameter "polyN". Besides the clustering idea, we add three simple proposal generation schemes based on the current solution as suggested in [4]. In *shift proposal*, the flow field in the current solution is shifted in either x or y directions for either 1, 2 or 3 pixels. In *stagger proposal*, the flow field is shifted by a vector randomly drawn from a Gaussian distribution. In *perturb proposal*, each flow value in the field is independently shifted by a vector randomly drawn from a Gaussian distribution. We choose schemes randomly when generating proposals.

4.3 Swarm Fusion layered depthmap estimation

Our last problem is layered depthmap estimation, recently proposed in [7] (see the anonymous paper in the supplementary material). The problem seeks to infer layered depthmap representation from a RGBD image, where each layer is a piecewise smooth segmented depthmap. This is essentially a multi-layer extension of Surface Stereo algorithm [15]. Layered depthmap estimation is a very challenging MRF inference problem due to its massive solution space. The number of labels per pixel is exponential in the number of layers, and is usually between 100,000 and 10,000,000. We copy their problem formulation and the proposal generation schemes. [1]

Competing methods

In this problem setting, solution proposals depend heavily on the current solution, eliminating the possibility of using Hierarchical Fusion Move (HFM), which needs to enumerate all the proposals to start. Therefore, viable competing methods are Fusion Move (FM) and Parallel Fusion Move (PFM) as in the optical flow problem. The fusions are binary, for which we use QPBO.

Swarm Fusion architectures

The three swarm architectures with the same configurations as in the optical flow problem have been evaluated.

5 Experimental results

We have implemented the algorithms with multi-threading support from C++ 11, and conducted the experiments on Linux PCs with Intel Core i7 4790 processor with 4 cores. We have used the Graph-cuts optimization code written by Veksler, using the libraries provided by Boykov and Kolmogorov [8, 22, 27, 28]. We have used the QPBO and TRW-S implementations by Kolmogorov [17, 23]. We have used 4 threads for experiments unless indicated. We now look at our experimental results for the three problems.

Stereo

We have chosen 7 images with the resolution of 695×555 from the Book sequence of Middlebury stereo dataset [24]. The number of disparity labels is set to 256. Since the order of labels is important for the expansion techniques, we have used the same random order for all algorithms to avoid any bias.

Figure 3 compares the converegence rate of the competing methods. Note that we define the energy of a multi-threading system to be the energy of the

[1] Authors have proposed a novel fusion scheme, where a solution subspace instead of a single solution is generated by a proposal generation scheme. Since a solution subspace can be represented by a concatenation of multiple solution proposals, their algorithm can be easily integrated into our Swarm Fusion framework. However, competing fusion methods (e.g., Parallel Fusion Move or Hierarchical Fusion) cannot handle a solution subspace proposal, making it impossible to conduct fair comparative evaluations. We choose to use a simple solution proposal for this experiment.

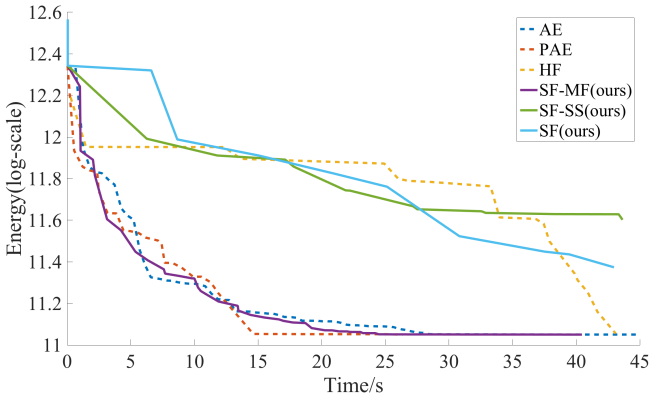

Fig. 3: Energy plots for the stereo problem. For a multi-threading algorithm, the energy at a given moment is defined to be the best solution so far.

best solution found so far. PAE and SF-MF converge faster than the single thread AE. However, the speedup is not significant, which confirms the fact that the problem is an easy one. Our approaches with multi-way fusion (SF-SS or SF) are the slowest kind, because the TRW-S for multi-way fusion is slower than multiple Alpha Expansion steps, and this stereo problem is too easy to gain benefits through mulit-way fusion. Figure 4 shows the energy plot per thread for SF-MF (ours) and PAE. With solution sharing, the energy in SF-MF decreases more uniformly, while in PAE the energy makes dramatic decrease at the final fusion.

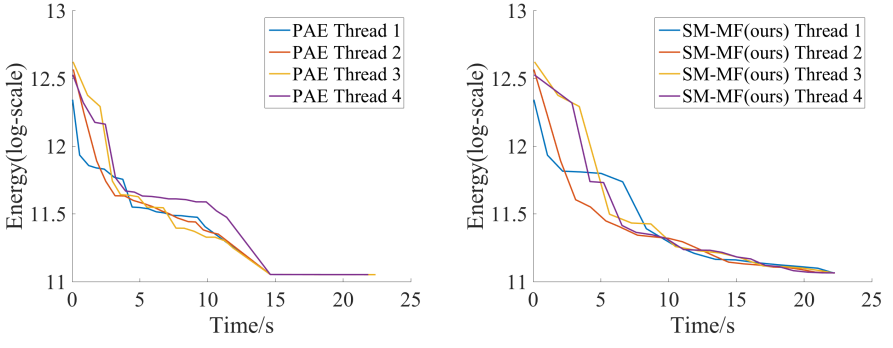

Fig. 4: Energy plots per thread for the stereo problem. Left: PAE. Right: SF-MF(ours).

For an easy optimization problem such as stereo with strong unary terms and submodular pairwise terms, our full architecture with solution sharing and multi-way fusion actually makes convergence slower compared with PAE due to its overhead.

Optical Flow

We have chosen the Dimetrodon image pair from the Middlebury flow dataset [25]. Figure 5 shows the energy plots of the three competing methods, Fusion Move (FM), Parallel Fusion Move (PFM), and Hierarchical Fusion Move (HFM), against our Swarm Fusion methods (SF-MF, SF-SS, SF). A key observation is that SF-MF converges quicker and better than PFM. This is indeed the benefits of solution sharing in our network. Optical flow is a more difficult problem and many solution proposals are not effective. The solution sharing (i.e., SF-MF) allows all the threads to exchange effective solution proposals in the middle of the optimization.

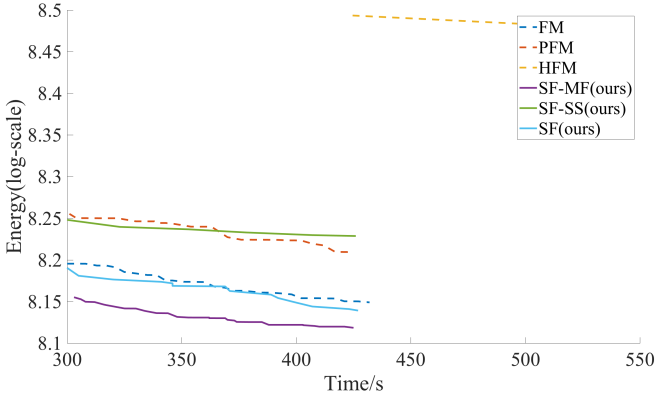

Fig. 5: Energy plots for the optical flow problem. SF-MF has the best performance due to its solution sharing strategy.

To further investigate the effectiveness of solution sharing, Figure 6 shows the energy plots of PFM and SF-MF per thread. As evident from the plot, in PFM, threads need to keep working independently at higher energy states. SF-MF, on the other hand, exchanges solutions all the time, and every thread is making an effective work in improving the solution. Another key finding from Fig. 5 is that SF is slower than SF-MF. Our analysis is that multi-way fusion is inefficient in this problem setting, since solution proposals are relatively independent and fusing the solution space would not gain much benefit. It rather loses performance against QPBO due to the overhead of TRW-S.

There are two factors influencing solution sharing: 1) *the number of solutions to share* and 2) *the frequency of solution sharing*. Both factors are controlled by

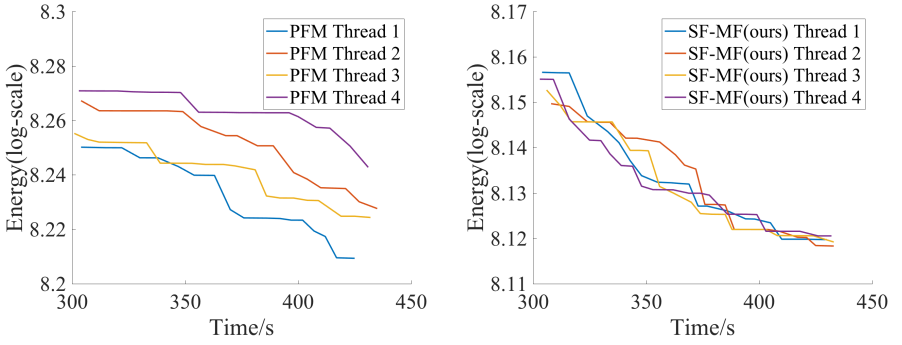

Fig. 6: Energy plots per thread for the optical flow problem. Left: Parallel Fusion Move (PFM). Right: SF-MF(ours).

β. As mentioned in Section 4.2, we have used $\beta = 1$ (i.e., share solutions) once in every five iterations. To further understand the effects of solution sharing, we conducted two more experiments. First, we set β to 0, 1, 2, or 3 in every five iterations, while keeping all other parameters the same (See Fig. 7(left)). Second, we change the number of iterations k between the two consecutive solution sharing iterations (See Fig. 7(right)). The first experiment revealed that the solution sharing makes convergence faster regardless of β. However, too much solution sharing slows down the convergence, and $\beta = 1$ is the sweet spot for this problem. The second experiment has shown that too frequent solution sharing harms the convergence, simply because threads have less time generating more proposals and exploring the solution space. Optimal parameter setting depends on each problem setting.

Layered depthmap estimation

We have used "ours_1" data in [7] for the experiments. Figure. 8 shows that Fusion Move, Parallel Fusion Move and SF-MF all got stuck in local minima, which is due to the lack of multi-way fusion. Layered depthmap estimation is a challenging problem with very large solution space. The binary fusion of solution proposals is too restrictive to make any improvements. This coincides with the observation in [7] that binary fusion of proposal solutions is not as powerful as their subspace fusion which is a special form of multi-way fusion here. Lastly, solution sharing also plays an important role for this challenging problem, as SF performs much better than SF-SS.

To further study the effects of multi-way fusion, we have varied the value of α which controls the number of solution proposals to be fused in SF-SS model (See Fig. 9(left)). Note that we have used SF-SS instead of SF to disable solution sharing and better observe the effects of multi-way fusion. It is interesting to see that more multi-way fusion takes longer to converge, but finds a lower energy state at the end.

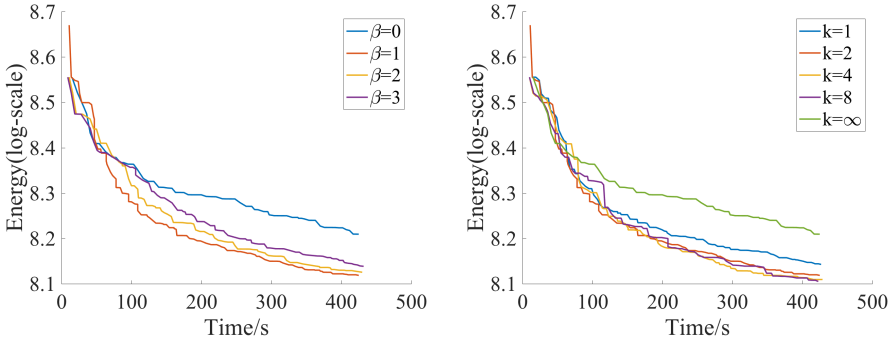

Fig. 7: Energy plots for optical flow under different configurations. Left: varying β. Right: varying solution sharing frequencies. Solution sharing achieves better convergence, but sharing too many solutions (larger β) or sharing solutions too frequently (less k) slows down the convergence, as it reduces the time for exploration.

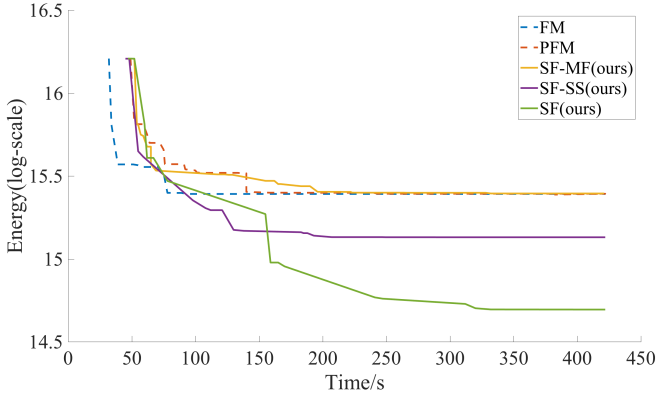

Fig. 8: Energy plots for the layered depthmap estimation problem. Both the multi-way fusion and the solution sharing are important for this challenging problem.

Finally, we have examined the role of multi-threading by varying the number of threads N in our most general model SF (See Fig. 9(right)) [1]. More threads lead to faster convergence as expected, although the rate of speed-up is not proportional to the number of threads due to the randomness in the proposal generation scheme.

[1]While keeping other parameters the same, we have to change β with N because of the constraint $\beta \leq N - 1$. We have always used $\beta = N - 1$ in this experiment.

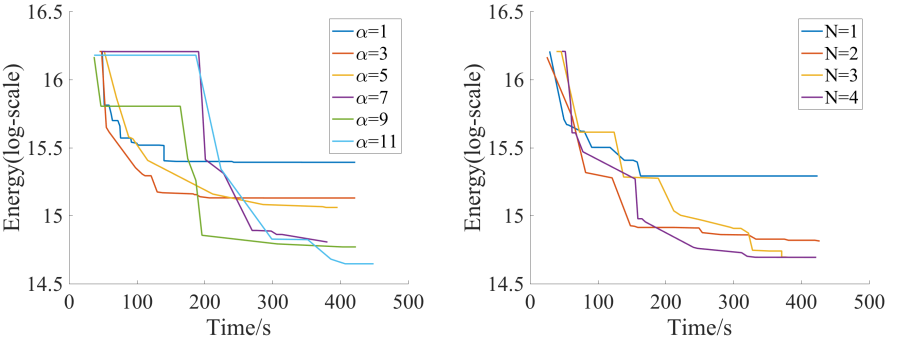

Fig. 9: Left: Energy plots for layered depthmap estimation with varying α. Right: Energy plots for layered depthmap estimation with varying number of threads N.

6 Conclusion and future directions

We have proposed a novel MRF inference framework, Swarm Fusion, in parallel computing environments. The framework is general and makes popular inference techniques such as Alpha Expansion, Fusion Move, Parallel Alpha Expansion, and Hierarchical Fusion, its special cases. Our experiments have revealed that the framework exploits parallel computational resources and achieves faster convergence, especially for challenging problems. Our first future work is to conduct experiments on cloud computing environments, in particular, the MapReduce programming model, where the roles of mappers and reducers exactly correspond to the processes of parallel multi-way fusion and solution sharing, respectively. Another future work is the automatic configuration of the Swarm Fusion architecture. Our experiments have shown that optimal architectures are different for different problems. An interesting direction is to adaptively change its architecture during the computation, for example, switching to simple parallel alpha-expansion for easy problems, or increasing the rate of solution exchanges when solutions vary significantly across threads. Parallel MRF inference has been a relatively under-explored topic in Computer Vision. The proposed Swarm Fusion framework can be intergrated into existing algorithms with minimal coding. We believe that this paper would immediately benefit tens of thousands of Computer Vision researchers or engineers in the world, who currently solve MRF problems. We will share our source code with the community.

7 Acknowledgement

This research was supported by National Science Foundation under grant IIS 1540012 and Google Faculty Research Award. We thank Nvidia for a generous GPU donation.

References

1. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al.: Mastering the game of go with deep neural networks and tree search. Nature **529**(7587) (2016) 484–489
2. Kopf, J., Uyttendaele, M., Deussen, O., Cohen, M.F.: Capturing and viewing gigapixel images. ACM Transactions on Graphics (TOG) **26**(3) (2007) 93
3. Hane, C., Zach, C., Cohen, A., Angst, R., Pollefeys, M.: Joint 3d scene reconstruction and class segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2013) 97–104
4. Lempitsky, V., Roth, S., Rother, C.: Fusionflow: Discrete-continuous optimization for optical flow estimation. In: Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on, IEEE (2008) 1–8
5. Lempitsky, V., Rother, C., Roth, S., Blake, A.: Fusion moves for markov random field optimization. Pattern Analysis and Machine Intelligence, IEEE Transactions on **32**(8) (2010) 1392–1405
6. Woodford, O., Torr, P., Reid, I., Fitzgibbon, A.: Global stereo reconstruction under second-order smoothness priors. Pattern Analysis and Machine Intelligence, IEEE Transactions on **31**(12) (2009) 2115–2128
7. Anonymous: Layered scene decomposition via the occlusion-crf
8. Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., Rother, C.: A comparative study of energy minimization methods for markov random fields with smoothness-based priors. Pattern Analysis and Machine Intelligence, IEEE Transactions on **30**(6) (2008) 1068–1080
9. Kappes, J.H., Andres, B., Hamprecht, F.A., Schnörr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B.X., Kröger, T., Lellmann, J., et al.: A comparative study of modern inference techniques for structured discrete energy minimization problems. International Journal of Computer Vision **115**(2) (2015) 155–184
10. Kumar, M.P., Koller, D.: Map estimation of semi-metric mrfs via hierarchical graph cuts. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, AUAI Press (2009) 313–320
11. Delong, A., Gorelick, L., Veksler, O., Boykov, Y.: Minimizing energies with hierarchical costs. International journal of computer vision **100**(1) (2012) 38–58
12. Veksler, O.: Efficient parallel optimization for potts energy with hierarchical fusion. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2015) 887–895
13. Batra, D., Kohli, P.: Making the right moves: Guiding alpha-expansion using local primal-dual gaps. In: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, IEEE (2011) 1865–1872
14. Strandmark, P., Kahl, F.: Parallel and distributed graph cuts by dual decomposition. In: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, IEEE (2010) 2085–2092
15. Bleyer, M., Rother, C., Kohli, P.: Surface stereo with soft segmentation. In: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on, IEEE (2010) 1570–1577
16. Beier, T., Hamprecht, F.A., Kappes, J.H.: Fusion moves for correlation clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2015) 3507–3516

17. Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. Pattern Analysis and Machine Intelligence, IEEE Transactions on **28**(10) (2006) 1568–1583
18. Murphy, K.P., Weiss, Y., Jordan, M.I.: Loopy belief propagation for approximate inference: An empirical study. In: Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence, Morgan Kaufmann Publishers Inc. (1999) 467–475
19. Man, K.F., Tang, K.S., Kwong, S.: Genetic algorithms: concepts and applications. IEEE transactions on Industrial Electronics **43**(5) (1996) 519–534
20. Kenndy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of IEEE International Conference on Neural Networks. Volume 4. (1995) 1942–1948
21. Sharp, T., Keskin, C., Robertson, D., Taylor, J., Shotton, J., Leichter, D.K.C.R.I., Wei, A.V.Y., Krupka, D.F.P.K.E., Fitzgibbon, A., Izadi, S.: Accurate, robust, and flexible real-time hand tracking
22. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. Pattern Analysis and Machine Intelligence, IEEE Transactions on **23**(11) (2001) 1222–1239
23. Rother, C., Kolmogorov, V., Lempitsky, V., Szummer, M.: Optimizing binary mrfs via extended roof duality. In: Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on, IEEE (2007) 1–8
24. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. International journal of computer vision **47**(1-3) (2002) 7–42
25. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. International Journal of Computer Vision **92**(1) (2011) 1–31
26. Furukawa, Y., Hernndez, C.: Multi-view stereo: A tutorial. Foundations and Trends in Computer Graphics and Vision **9**(1-2) (2013) 1–148
27. Kolmogorov, V., Zabin, R.: What energy functions can be minimized via graph cuts? Pattern Analysis and Machine Intelligence, IEEE Transactions on **26**(2) (2004) 147–159
28. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. In: Energy minimization methods in computer vision and pattern recognition, Springer (2001) 359–374